

Inferring interaction partners from protein sequences using mutual information

Anne-Florence Bitbol^{1*}

1 Sorbonne Université, CNRS, Laboratoire Jean Perrin (UMR 8237), F-75005 Paris, France

* anne-florence.bitbol@sorbonne-universite.fr

Abstract

Specific protein-protein interactions are crucial in most cellular processes. They enable multi-protein complexes to assemble and to remain stable, and they allow signal transduction in various pathways. Functional interactions between proteins result in coevolution between the interacting partners, and thus in correlations between their sequences. Pairwise maximum-entropy based models have enabled successful inference of pairs of amino-acid residues that are in contact in the three-dimensional structure of multi-protein complexes, starting from the correlations in the sequence data of known interaction partners. Recently, algorithms inspired by these methods have been developed to identify which proteins are specific interaction partners among the paralogous proteins of two families, starting from sequence data alone. Here, we demonstrate that a slightly higher performance for partner identification can be reached by an approximate maximization of the mutual information between the sequence alignments of the two protein families. This stands in contrast with structure prediction of proteins and of multi-protein complexes from sequence data, where pairwise maximum-entropy based global statistical models substantially improve performance compared to mutual information. Our findings entail that the statistical dependences allowing interaction partner prediction from sequence data are not restricted to the residue pairs that are in direct contact at the interface between the partner proteins.

Author summary

Specific protein-protein interactions are at the heart of most intra-cellular processes. Mapping these interactions is thus crucial to a systems-level understanding of cells, and has broad applications to areas such as drug targeting. Systematic experimental identification of protein interaction partners is still challenging. However, a large and rapidly growing amount of sequence data is now available. Recently, algorithms have been proposed to identify which proteins interact from their sequences alone, thanks to the co-variation of the sequences of interacting proteins. These algorithms build upon inference methods that have been used with success to predict the three-dimensional structures of proteins and multi-protein complexes, and their focus is on the amino-acid residues that are in direct contact. Here, we propose a simpler method to identify which proteins interact among the paralogous proteins of two families, starting from their sequences alone. Our method relies on an approximate maximization of mutual information between the sequences of the two families, without specifically emphasizing the contacting residue pairs. We demonstrate that this method slightly outperforms the earlier one. This result highlights that partner prediction does not only rely on the identities and interactions of directly contacting amino-acids.

Keywords

Protein-protein interactions; mutual information; inference; sequence data; information theory; paralogs; homologs

Introduction

Most cellular processes are carried out by interacting proteins. Specific protein-protein interactions allow multi-protein complexes to assemble, and ensure proper signal transduction in various pathways. Hence, mapping specific protein-protein interactions is an important fundamental question. High-throughput experiments have recently elucidated a substantial fraction of protein-protein interactions in a few model organisms [1], but such experiments remain challenging. An attractive alternative possibility is to exploit the increasingly abundant sequence data in order to identify specific protein-protein interaction partners.

The sequences of interacting proteins are correlated, both because of evolutionary constraints arising from the need to maintain physico-chemical complementarity among amino-acids in contact, and because of shared evolutionary history. The first type of correlations has received substantial interest, both within single proteins and across protein partners, as evolutionary constraints induce correlations between amino acids that are in contact in the folded protein or in the multi-protein complex. Hence, the correlations observed in multiple sequence alignments of homologous proteins contain information about protein structure. Global statistical models allow direct and indirect correlations to be disentangled [2–4]. Such models, built using the maximum entropy principle [5], and assuming pairwise interactions, known in the field of proteins as Direct Coupling Analysis (DCA), have been used with success to determine three-dimensional protein structures from sequences [6, 7], to predict mutational effects [8–10], to find residue contacts between known interaction partners [4, 11–15], and most recently to predict interaction partners from sequence data [16, 17]. DCA models lay the emphasis on interactions between residues that are in direct contact in the three-dimensional protein structure, and have been optimized for contact prediction. However, correlations in protein sequences also have important collective modes [18, 19], which can arise from functional selection [20, 21], and additional correlations are due to phylogeny [18, 22, 23]. These contributions are deleterious to the prediction of contacts [23] but not necessarily to the prediction of interacting partners, since a pair of interacting partners may be subject to common functional selection, and may also have a more strongly shared phylogenetic history than non-interacting proteins.

Here, we present an alternative approach to predict interaction partners from sequence data, among the paralogous proteins belonging to two interacting families. In contrast to the previous pairwise maximum entropy-based approaches [16, 17], the present method is based on an approximate maximization of mutual information between the sequences from the two protein families. Specifically, we develop a variant of the iterative pairing algorithm (IPA) introduced in Ref. [16], where we use mutual information (MI) as a score to maximize, instead of the effective interaction energy from a pairwise maximum entropy (DCA) model. We demonstrate that this mutual information-based algorithm (MI-IPA) performs slightly better than the one (DCA-IPA) introduced by us and colleagues in Ref. [16]. Our findings entail that the statistical dependences allowing interaction partner prediction from sequence data are not restricted to the contacting residue pairs revealed by DCA.

Results

We developed an iterative pairing algorithm (MI-IPA) that pairs paralogous proteins from two interacting protein families \mathcal{A} and \mathcal{B} by approximately maximizing mutual information between the sequences of the two families. Here, we first introduce the information theory-based pairing score we employ, before briefly explaining the steps of the MI-IPA. Next, we present the results we obtained with the MI-IPA. Throughout, we compare the performance of the MI-IPA to that obtained with the DCA-IPA from Ref. [16], which infers a pairwise maximum entropy model and approximately maximizes the resulting effective interaction energies. First, we consider the case where the MI-IPA starts with a training set of known protein pairs, obtaining good performance even with few training pairs. Then, we show that the MI-IPA can make accurate predictions starting without any training set, as would be needed to predict novel protein-protein

interactions. We further demonstrate the robustness of our algorithm by successfully applying it to several pairs of proteins. Finally, we assess to what extent the MI-IPA is successful at maximizing mutual information.

A pairing score based on pointwise mutual information (PMI)

Consider an alignment of M concatenated sequences AB of length L , where A is a protein from family \mathcal{A} and B is a protein from family \mathcal{B} . At each amino-acid site $i \in \{1, \dots, L\}$, a given sequence can feature any amino acid (represented by $\alpha \in \{1, \dots, 20\}$), or a gap (represented by $\alpha = 21$). To describe the statistics of this concatenated alignment (CA), we employ the single-site frequencies of occurrence of each state α at each site i , denoted by $f_i(\alpha)$, and the two-site frequencies of occurrence of each ordered pair of states (α, β) at each ordered pair of sites (i, j) , denoted by $f_{ij}(\alpha, \beta)$. These empirical frequencies are obtained by counting the sequences where given residues occur at given sites and dividing by the number M of sequences in the CA. (Note that when computing frequencies from real protein data, it is useful to weight sequences so as to attenuate the impact of biased sampling, and to include pseudocounts, in order to mitigate finite-sample effects, see Methods.) The empirical frequencies constitute estimates of the corresponding probabilities $p_i(\alpha)$ and $p_{ij}(\alpha, \beta)$, and tend toward them in the limit where the number M of sequences in the alignment tends to infinity.

The pointwise mutual information (PMI) of a pair of residues (α, β) at a pair of sites (i, j) is defined as [24]:

$$\text{PMI}_{ij}(\alpha, \beta) = \log \left[\frac{p_{ij}(\alpha, \beta)}{p_i(\alpha)p_j(\beta)} \right]. \quad (1)$$

Averaging this quantity over all possible residue pairs yields the mutual information (MI) between sites i and j [25]:

$$\text{MI}_{ij} = \sum_{\alpha, \beta} p_{ij}(\alpha, \beta) \text{PMI}_{ij}(\alpha, \beta) = \sum_{\alpha, \beta} p_{ij}(\alpha, \beta) \log \left[\frac{p_{ij}(\alpha, \beta)}{p_i(\alpha)p_j(\beta)} \right]. \quad (2)$$

PMI has been used in linguistics to study the co-occurrence of words [26, 27]. Note that in some instances [24, 26] PMI is called MI, and MI is then referred to as the average value of MI.

We define a pairing score S_{AB} for each pair AB of proteins as the sum of the PMIs of the inter-protein pairs of sites of this concatenated sequence (i.e. those that involve one site in protein A and one site in protein B):

$$S_{AB} = \sum_{i=1}^{L_A} \sum_{j=L_A+1}^L \text{PMI}_{ij}(\alpha_i, \beta_j), \quad (3)$$

where we have denoted the concatenated sequence AB by $(\alpha_1, \dots, \alpha_{L_A}, \beta_{L_A+1}, \dots, \beta_L)$, with L_A the length of the A sequence. This score can be computed for a pair AB that is a member of the CA used to estimate the PMI of each residue pair at each site, but also for any other pair AB comprised of the sequences of members of the protein families \mathcal{A} and \mathcal{B} .

Next, consider a candidate assignment X of M' pairs AB, where each protein A is paired with a protein B from the same species, resulting in a CA of M' sequences of length $L = L_A + L_B$. Again, the pairs in this CA can involve the proteins in the CA of M pairs used to estimate the PMIs, with the same assignment or a different one, or any other pair AB comprised of the sequences of members of the protein families \mathcal{A} and \mathcal{B} . We define the overall pairing score S_X of the assignment X by the average of all pairing scores (see Eq. 3) of the pairs involved:

$$\begin{aligned} S_X &= \frac{1}{M'} \sum_{AB \in X} S_{AB} = \frac{1}{M'} \sum_{i=1}^{L_A} \sum_{j=L_A+1}^L \sum_{AB \in X} \text{PMI}_{ij}(\alpha_i, \beta_j) \\ &= \frac{1}{M'} \sum_{i=1}^{L_A} \sum_{j=L_A+1}^L \sum_{\alpha, \beta} \sum_{\substack{AB \in X \\ \alpha_i = \alpha, \beta_j = \beta}} \text{PMI}_{ij}(\alpha, \beta) = \sum_{i=1}^{L_A} \sum_{j=L_A+1}^L \sum_{\alpha, \beta} f_{ij}(\alpha, \beta) \text{PMI}_{ij}(\alpha, \beta). \end{aligned} \quad (4)$$

In the limit of large alignments, the empirical frequencies tend toward probabilities. Besides, if the scored CA is the same as the one used to calculate the PMIs, all frequencies will be the same for both of them. For different CA, in the case of an assignment consistent with the CA used to estimate the PMIs, the limiting two-body probabilities will be the same in the two CA. Hence, in these cases, combining Eqs. 2 and 4 yields

$$S_X \xrightarrow[\substack{M \rightarrow \infty \\ M' \rightarrow \infty}]{\substack{L_A \\ L}} \sum_{i=1}^{L_A} \sum_{j=L_A+1}^L MI_{ij}. \quad (5)$$

For large alignments, maximizing S_X thus corresponds to maximizing the sum of the MIs of inter-protein site pairs, which is itself a pairwise approximation of the MI between the sequences of two protein families \mathcal{A} and \mathcal{B} . A brute-force maximization of S_X over all possible assignments X of a realistic dataset would result in a combinatorial explosion, since all allowed permutations of pairs would need to be considered. In practice, since biologically meaningful pairings have to be made within a species, this means that we would need to consider all combinations of all permutations within each species, which already yields prohibitively large numbers of assignments to test. Hence, we propose an algorithm to perform an approximate maximization of S_X .

An iterative pairing algorithm (IPA) based on MI

In order to approximately maximize mutual information via the score S_X (see Eq. 4), we propose an iterative pairing algorithm (referred to as MI-IPA) inspired by that of Ref. [16] (referred to as DCA-IPA), where the effective interaction energy from a global statistical model was approximately maximized.

In each iteration, we first estimate PMIs for all inter-protein residue pairs from a concatenated alignment (CA) of paired sequences. The initial CA, used at the first iteration, is either built from a training set of known correct protein pairs, or made from random pairs, assuming no prior knowledge of interacting pairs. We calculate the pairing scores S_{AB} (see Eq. 3) for every possible protein pair AB within each species, by summing the inter-protein PMIs. Next, within each species, we assign pairs by maximizing the sum of S_{AB} scores in the species (assuming one-to-one specific interactions), thereby maximizing S_X (see Eq. 4) over biologically relevant pair assignments, where each protein has a partner within its species. We attribute a confidence score to each predicted pair, by using the difference of scores between the optimal assignment of pairs in the species and the best alternative assignment that does not involve this predicted pair. The CA is then updated by including the highest-scoring protein pairs, and the next iteration can begin. At each iteration, all pairs in the CA are re-selected based on confidence scores (except the initial training pairs, if any), allowing for error correction. More details on each step of the MI-IPA are given in Methods.

The MI-IPA accurately predicts interaction partners from a training set of known partners

As in Ref. [16], we use histidine kinases (HKs) and response regulators (RRs) from prokaryotic two-component signaling systems as our main benchmark. Two-component systems are important pathways that enable bacteria to sense and respond to environment signals. Typically, a transmembrane HK senses a signal, autophosphorylates, and transfers its phosphate group to its cognate RR, which in turn induces a cellular response [28]. Importantly, most cognate HK-RR pairs are encoded in the same operon, so actual interaction partners are known, which enables us to assess performance.

Unless otherwise specified, our results were obtained on a “standard dataset” comprising 5064 HK-RR pairs for which the correct pairings are known from gene adjacency. Each species has on average $\langle m_p \rangle = 11.0$ pairs, and at least two pairs (see Methods).

We start by predicting interaction partners starting from a training set of known pairs. As our training set, we pick a random set of N_{start} known HK-RR pairs from the standard dataset.

The first iteration of the MI-IPA uses this concatenated alignment (CA) to compute PMIs and score possible pairs. We blind the pairings of the remaining dataset, and use it as a testing set on which we predict pairings. At each subsequent iteration $n > 1$, the CA used to recompute PMIs contains the initial training pairs plus the $(n - 1)N_{\text{increment}}$ highest-scoring predicted pairs from the previous iteration (see Methods).

As in the case of the DCA-IPA [16], iterating, and thereby progressively adding high-scoring pairs to the CA, allows us to increase the fraction of pairs that are correctly predicted. This gradual improvement of the TP fraction during the iterations of the MI-IPA is shown in Fig. 1A for different training set sizes N_{start} . The increase of TP fraction is especially spectacular for small training sets. Fig. 1B shows the initial TP fraction, obtained at the first iteration, and the final TP fraction, obtained at the last iteration, versus the size of the training set N_{start} , both for the MI-IPA and for the DCA-IPA [16]. In both cases, comparing the initial and final TP fractions demonstrates the major interest of our iterative approach, through the massive increase in TP fraction, especially for small training sets. Moreover, the final TP fraction depends only weakly on N_{start} : the iterative approach removes the need for large training sets. Both algorithms yield very good performance, and the MI-IPA even outperforms the DCA-IPA in the trickiest case of small training sets. In this limit ($N_{\text{start}} = 1$), the MI-IPA yields 86% true positive (TP) pairs while the DCA-IPA yields 84% TP, while both start from 12% TP at the first iteration. These final TP fractions are strikingly higher than the random expectation of 9%, while such small training sets contain very little information about pairings, as illustrated by the associated low initial TP fraction.

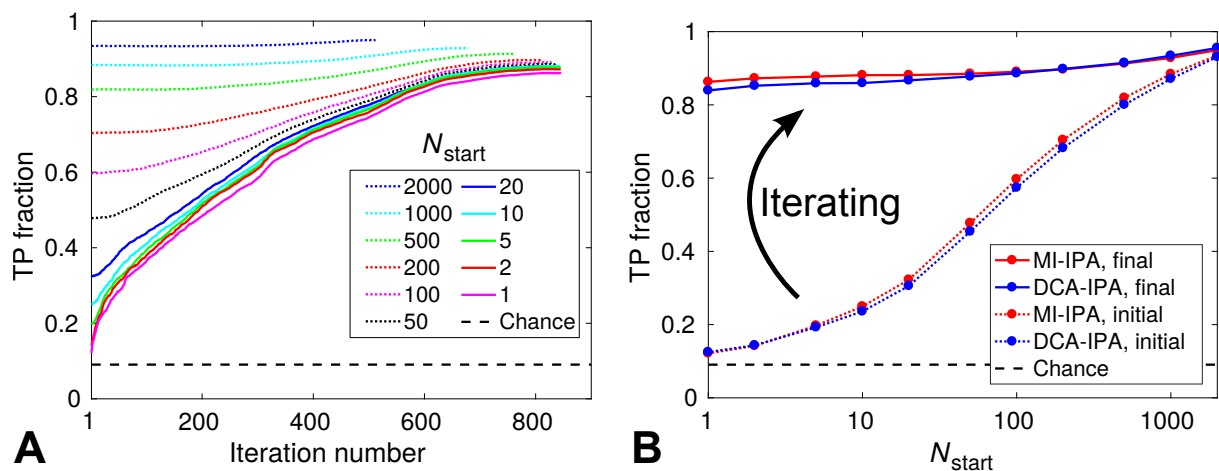


Figure 1: **Performance of the MI-IPA for different training set sizes N_{start} .** (A) Increase of the TP fraction with the number of iterations during the MI-IPA. (B) Initial and final TP fractions (at the first and last iteration) versus N_{start} for the MI-IPA and for the DCA-IPA of Ref. [16]. In both panels, the standard dataset of HK-RRs is used, and the CA includes $N_{\text{increment}} = 6$ additional pairs at each iteration. All results are averaged over 50 replicates that differ by the random choice of HK-RR pairs in the training set. Dashed lines represent the average TP fraction obtained for random within-species HK-RR pairings.

Since the MI-IPA to the DCA-IPA yield similar performance, we asked whether they tend to predict the same correct pairs when starting from the same training set. To assess this, consider a species with m AB pairs, and denote by p (resp. q) the number of pairs correctly assigned by the MI-IPA (resp. by the DCA-IPA) in this species. If the two algorithms made independent predictions, it would correspond to independently and randomly drawing p (resp. q) proteins A among m , to be correctly paired. In this null model, the number of possible MI-IPA assignments that share k correct pairs with the DCA-IPA is $\binom{q}{k} \binom{m-q}{p-k}$: k correct pairs are chosen among the q pairs correctly assigned by the DCA-IPA, and the other $p - k$ ones are chosen among the $m - q$ proteins A incorrectly paired by the DCA-IPA. Besides, the total number of possible MI-IPA assignments with p correct pairs is $\binom{m}{p}$. Hence, the probability $P(p, k, m, q)$ that k correct pairs

are predicted by both algorithms is given by the hypergeometric distribution:

$$P(p, k, m, q) = \frac{\binom{q}{k} \binom{m-q}{p-k}}{\binom{m}{p}}. \quad (6)$$

The expectation of k under this distribution is given by

$$\langle k \rangle = \frac{pq}{m}. \quad (7)$$

Hence, in this fully independent null model, the expectation of the total number of correct pairs predicted by both algorithms can be obtained by summing the expectations in Eq. 7 over all species in the dataset. Moreover, in each species, the observed number k_{obs} of shared correct pairs can be compared to this null model, as well as to the extreme case where all correct pairs that can be shared are shared, yielding $\min(p, q)$ shared pairs. Hence, we define the relative excess E of shared predictions by

$$E = \frac{\sum_{i=1}^S k_{\text{obs},i} - \langle k_i \rangle}{\sum_{i=1}^S \min(p_i, q_i) - \langle k_i \rangle}, \quad (8)$$

where the index i corresponds to a particular species, the sums run over the S species present in the dataset, and the expectations are given by Eq. 7. If $E > 0$, the two algorithms tend to predict the same pairs more frequently than if their predictions were fully independent. In addition, the maximal value E can take is 1, including the case where predictions from both algorithms are exactly the same. We calculated the average value of E across 50 replicates where both algorithms were started from the same training set, for various N_{start} values, in the same conditions as in Fig. 1. As expected, we found that E increases with N_{start} , as the algorithms share more information to begin with. However, E depends rather weakly on the size N_{start} of the training set, varying smoothly from 53% for $N_{\text{start}} = 1$ to 68% for $N_{\text{start}} = 2000$. This indicates a significant tendency of the two algorithms to make the same correct predictions, even in the case of small training sets.

Great accuracy is maintained in the absence of a training set

In Ref. [16], we showed that the DCA-IPA yields very good identification of interacting pairs without any training set, i.e. without any prior knowledge of interacting pairs. Given this previous result, and given the success of the MI-IPA with very small training sets, we ask whether the MI-IPA also makes good predictions in the absence of a training set. To test this, we followed the approach introduced in Ref. [16] by randomly pairing each HK with an RR from the same species, and using these 5064 random pairs to train the initial model. At each subsequent iteration $n > 1$, the CA only contained the $(n - 1)N_{\text{increment}}$ highest-scoring pairs from the previous iteration (see Methods).

Fig. 2A shows the final TP fraction obtained for different values of $N_{\text{increment}}$, both for the MI-IPA and for the DCA-IPA. In both cases, the iterative method performs best for small increment steps, which highlights again the interest of the iterative approach. Importantly, the MI-IPA performs better than the DCA-IPA for all values of $N_{\text{increment}}$, and requires substantially less small increments than the DCA-IPA to reach the same performance. The low- $N_{\text{increment}}$ limit of the final TP fraction is 0.87 for the MI-IPA, versus 0.84 for the DCA-IPA. These values are consistent with those obtained above with a single training pair, $N_{\text{start}} = 1$ (Fig. 1A). We emphasize that the striking TP fraction of 0.87 is attained by the MI-IPA without any prior knowledge of HK-RR interactions. Ref. [16] showed that an important ingredient for the DCA-IPA to bootstrap its way toward high predictive power is that sequence similarity is favored at early iterations, which increases the TP fraction in the CA, because correct pairs have more neighbors in terms of Hamming distance than incorrect pairs. In Ref. [16], this was called the *Anna Karenina effect*, in reference to the first sentence of Tolstoy's novel. The same explanation holds for the success of the MI-IPA starting from no training set.

In Fig. S1A, we investigated the performance of other variants of the IPA in the same conditions as in Fig. 2A. Instead of PMIs, we explored the possibilities of using normalized PMIs (NPMIs), $\text{NPMI}_{ij}(\alpha, \beta) = -\text{PMI}_{ij}(\alpha, \beta) / \log(f_{ij}(\alpha, \beta))$ [27] and covariances, $C_{ij}(\alpha, \beta) = f_{ij}(\alpha, \beta) - f_i(\alpha)f_j(\beta)$ (recall that in the mean-field approximation, DCA employs the inverse of this covariance matrix). For small $N_{\text{increment}}$ values, we find that NPMIs perform similarly as PMIs, while covariances do significantly worse than both PMIs and DCA scores.

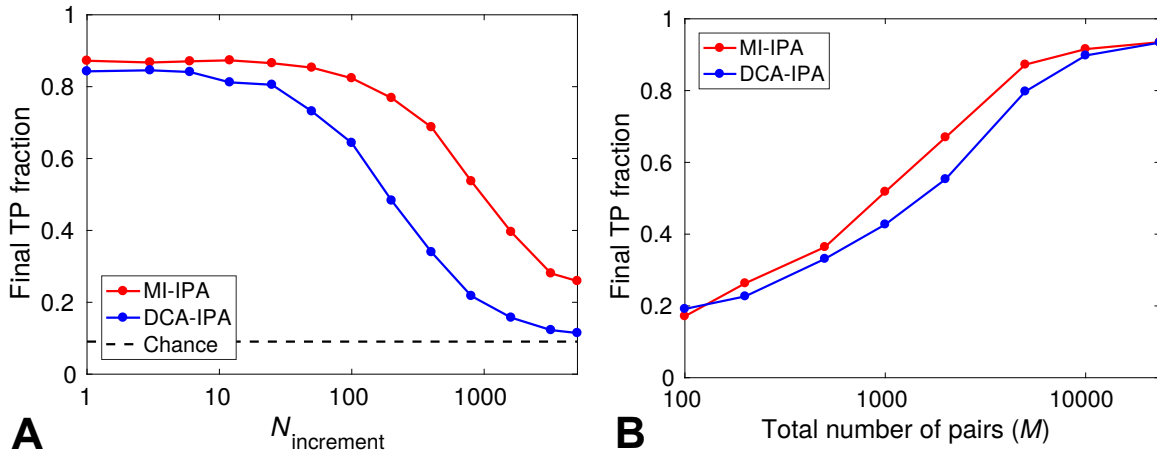


Figure 2: Performance of the MI-IPA without a training set. (A) Final TP fraction obtained by the MI-IPA and the DCA-IPA versus $N_{\text{increment}}$ for the standard HK-RR dataset. At the first iteration, the CA is made of random within-species HK-RR pairs. At each subsequent iteration $n > 1$, the CA includes the $(n-1)N_{\text{increment}}$ top predicted pairs. All results are averaged over 50 replicates employing different initial random pairings. The dashed line represents the average TP fraction obtained for random within-species HK-RR pairings. (B) Final TP fraction obtained by the MI-IPA and the DCA-IPA versus the total number M of HK-RR pairs in the dataset, starting from random pairings. For each M , except that corresponding to the full dataset, datasets are constructed by picking species randomly from the full dataset, and results are averaged over multiple different such alignments (from 50 up to 500 for small M). For the full dataset (largest M), averaging is done on 50 different initial random within-species pairings. All results in (B) are obtained in the small- $N_{\text{increment}}$ limit.

An important parameter for the performance of the MI-IPA is dataset size. Qualitatively, larger datasets imply more close neighbors, which is favorable to the success of bootstrapping, and they also allow one to estimate MI more accurately, so we expect the MI-IPA to perform best for large datasets. In the case of DCA, it is well-known that large datasets are required to properly infer the global statistical model at the heart of the method, so it is interesting to study whether the two methods will have similar limitations for small datasets or not. Indeed, Fig. 2B shows that the performance of both methods increases with dataset size. However, the rise of performance occurs for slightly smaller datasets in the case of the MI-IPA than for the DCA-IPA. Hence, the MI-IPA is better suited for partner prediction in small datasets. For the complete dataset (23,424 HK-RR pairs, see Methods), both methods reach the same striking final TP fraction of 0.93.

As above, we investigated the extent to which the MI-IPA and the DCA-IPA tend to make the same correct predictions. Note that this time, the two algorithms do not start from shared information since there is no training set. For the standard dataset of 5064 HK-RR pairs, with $N_{\text{increment}} = 6$, we found that the average excess shared fraction (see Eq. 7) of correct pairs between the two algorithms is $E = 53\%$, i.e. the same as the value obtained in the limit of a very small shared training set, $N_{\text{start}} = 1$ (see above). Recall that this positive value means that the two algorithms tend to make the same correct predictions. Moreover, the bias is 53% of the maximal value it could take. It is interesting to compare this result to the bias toward shared correct predictions across different replicates of the same algorithm. We obtained $E = 73\%$ for

the DCA-IPA, and $E = 94\%$ for the MI-IPA, meaning that predictions from the two different algorithms are less similar than those made by the same algorithm.

Since the two algorithms do make some different correct predictions, we next asked whether this can be exploited, with the intuition that pairs that are predicted both by the MI score and by the DCA score will tend to be correct more often than other pairs. We tried a combined IPA which calculates both scores at each iteration and computes separately the two corresponding pair assignments. Our ranking of pairs puts first the pairs contained in both assignments, and these pairs are ordered by decreasing MI-based confidence scores. Next come the other pairs: the assignment predicted using MI is conserved for them, and they are ordered by decreasing MI-based confidence scores. This means that the pairs consistently predicted by both methods are going to enter the CA earlier in the iterative process than those that differ. A very minor improvement was obtained over the MI-IPA using this method (see Fig. S1B). Since DCA typically requires large datasets to be reliable, we also tried combining MI and DCA by using MI at early iterations, when the CA is small, and switching to DCA at later iterations. No improvement over the MI-IPA was obtained using this method (see Fig. S1B).

Performance of the MI-IPA is robust across various protein pairs

To demonstrate the generality of the MI-IPA, we applied it to several pairs of protein families [13]. First, we considered several protein pairs involved in ABC transporter complexes, which permit the translocation of different substances across cell membranes [29]. We built alignments of homologs of the *Escherichia coli* interacting protein pairs MALG-MALK, FBPB-FBPC, and GSIC-GSID, all involved in ABC transporter complexes, using the same method as in Ref. [16] (see Methods). As in the case of HK-RRs, for each of these pairs of protein families, we worked on subsets of ~ 5000 protein pairs from species containing at least two pairs. In addition, we considered smaller families of proteins, yielding ~ 2000 pairs. More specifically, we chose two pairs of proteins involved in enzymatic complexes: PAAH-PAAJ is a pair of proteins involved in the fatty acid β -oxidation multienzyme complex [30], and XDHC-XDHA is a pair of proteins involved in the xanthine dehydrogenase complex [31].

Fig. 3 shows the performance of the MI-IPA and of the DCA-IPA for these protein pairs, starting from no initial training set, for various values of $N_{\text{increment}}$. In all cases, we find that both methods perform very well for small $N_{\text{increment}}$, but that higher performances are reached by the MI-IPA, particularly at larger $N_{\text{increment}}$ values. This is consistent with the results obtained with HK-RRs (see Fig. 2A). Note that here, compared to HK-RRs, larger $N_{\text{increment}}$ values suffice to obtain good performance. This is due to the fact that the pairing tasks are easier here, since the average number of proteins pairs per species is smaller (see Fig. 3, to be compared to $\langle m_p \rangle = 11.0$ for HK-RRs).

These accurate predictions demonstrate the broad applicability of the MI-IPA. Note that HK-RRs interact transiently, while the ABC transporter and enzymatic proteins we considered form permanent complexes, which highlights the generality of the method.

The MI-IPA reaches near-maximal MI

The MI-IPA approximately maximizes MI between the sequence alignments of two protein families \mathcal{A} and \mathcal{B} . However, there is not guarantee that it will find the assignment with highest MI. In addition, the score we maximize converges toward the sum of MIs of all inter-protein residue pairs (henceforth called “pairwise MI”) only in the limit of large alignments and assignments consistent with the CA used to estimate the PMIs (see Eq. 5). In practice, how well does the MI-IPA approach the goal of maximizing pairwise MI?

In order to answer this question, we now compare the pairwise MI of the pairs assigned by the MI-IPA to that of the actual protein pairs, and to the pairwise MI of the random within-species assignment which is used to initialize the MI-IPA in the absence of a training set. Fig. 4A shows these three quantities as a function of dataset size M , for HK-RR datasets to which the MI-IPA is applied starting from no training set, as in Fig. 2B. We observe a global trend of all computed

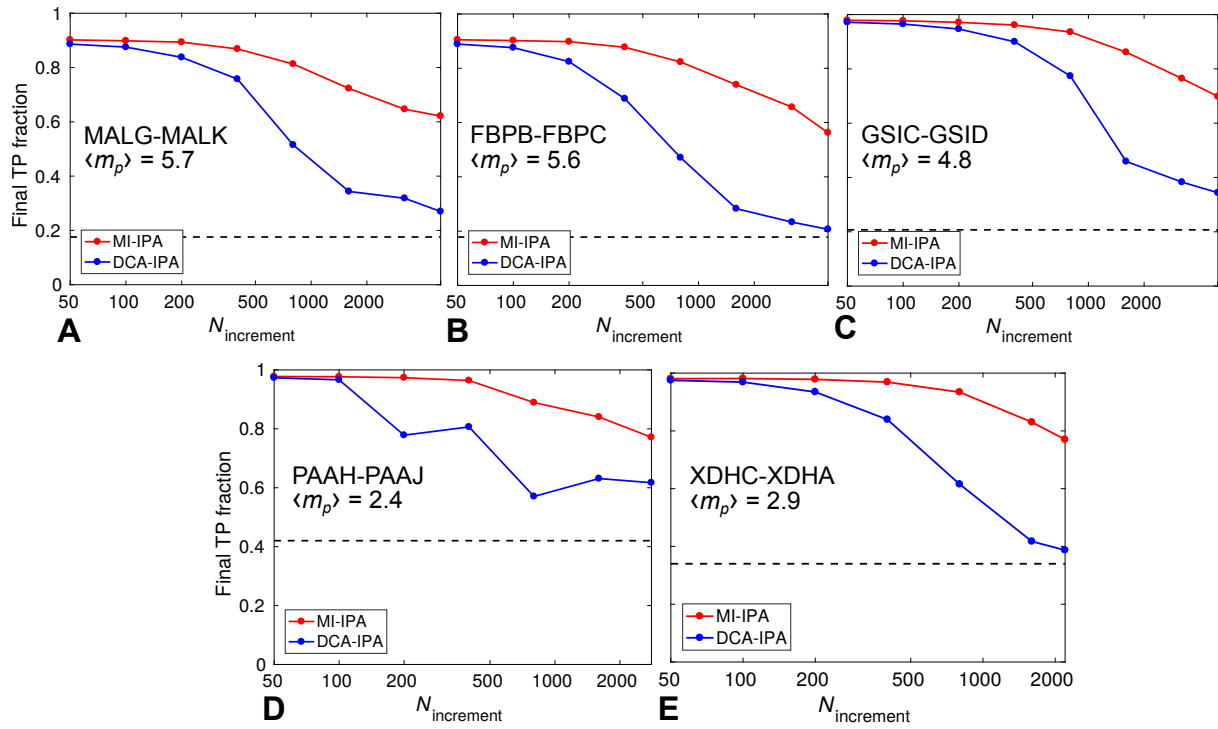


Figure 3: Performance of the MI-IPA for different protein pairs. Final TP fraction obtained without a training set by the MI-IPA and the DCA-IPA versus $N_{\text{increment}}$. (A,B,C) Pairs involved in ABC transporters; datasets of ~ 5000 pairs extracted from larger full datasets. (D,E) Pairs involved in enzymatic complexes; full datasets of ~ 2000 pairs. In each case, the mean number $\langle m_p \rangle$ of pairs per species is indicated, and dashed lines represent the average TP fraction obtained for random within-species pairings. All results are averaged over 50 replicates that differ in their initial random within-species pairings.

pairwise MIs to decrease when M is increased. This arises from a well-known finite size effect that occurs when estimating entropies from real datasets [32], and thus affects entropy-derived quantities such as MIs. To illustrate this point, Fig. 4A also shows the pairwise MI of HK-RR datasets where each column of the alignment is randomly scrambled, thus destroying actual correlations while retaining finite-size noise, as well as one-body frequencies. This null model features a similar decreasing trend as the other curves, thus demonstrating that this trend comes from finite-size effects. Note also that the scrambled alignment features pairwise MI values close to those of the initial random assignment, which makes sense because in both cases inter-protein residue pairs are decorrelated by the scrambling. The slightly higher MI of the initial random assignment arises from the fact that random partners are chosen within each species in this case, while complete columns are scrambled in the null model. Apart from the downward trend, a striking observation from Fig. 4A is that the pairwise MIs of the assignments predicted by the MI-IPA are significantly higher than those of the initial random assignments, and close to those of the actual protein pairs. Fig. 4B highlights these points by considering the excess pairwise MI in the actual protein pairs versus the initial and final assignments in the MI-IPA. Interestingly, even for small datasets, where the MI-IPA yields small TP fractions (see Fig. 2B), the pairwise MI of the assignment predicted by the MI-IPA is much closer to that of the actual assignment than to that of the initial random assignment. This suggests that the MI-IPA does a good job at maximizing pairwise MI, even though it does not reach the absolute maximum.

So far, we have used naive estimates of the pairwise MI, employing empirical frequencies instead of probabilities in MI (see Eq. 2), without correcting for the finite-size effect discussed above and visible in Fig. 4A. Various approaches have been proposed in order to correct for the systematic error due to finite-size effects in entropy (and thus MI) estimates. One can use the fact that these finite-size effects have a leading-order contribution of the form $\widehat{\text{MI}} - \text{MI} \propto K/M$, where

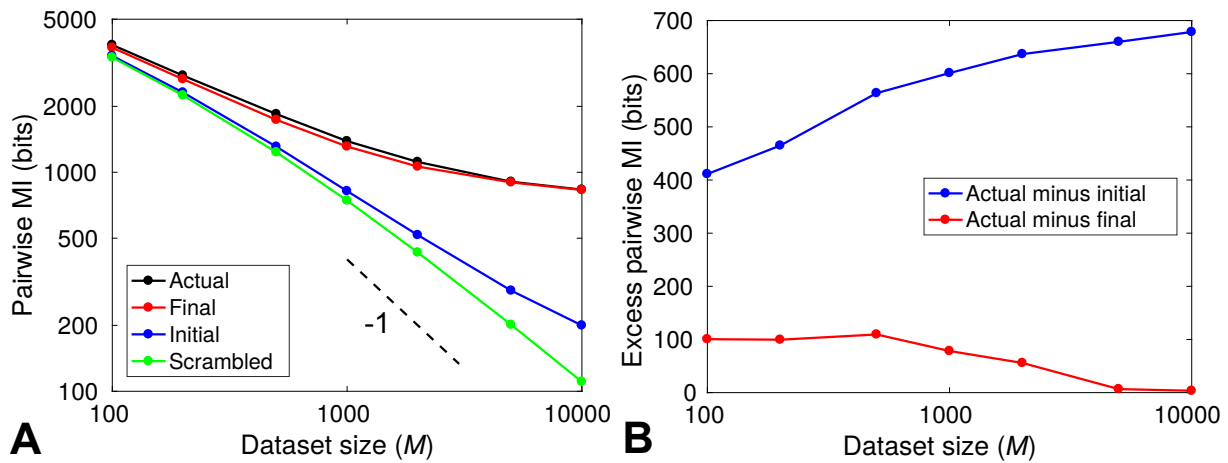


Figure 4: **Increase of pairwise MI obtained by the MI-IPA.** (A) Estimates of the pairwise MI (sum of MIs of all inter-protein residue pairs) are shown versus the total number M of sequences in the dataset. The different curves correspond to the pairwise MI of the actual set of correctly paired HK-RRs, of the initial random within-species assignment used to initialize the MI-IPA, and of the final assignment predicted by the MI-IPA, as well as of an alignment where each column was scrambled. Both axes have a logarithmic scale. The slope -1 , expected from leading-order finite-size effects, is indicated by the dashed line. For each M , HK-RR datasets are constructed by picking species randomly from the full dataset, and results are averaged over 50 different such alignments. (B) Excess pairwise MI of the actual set of correctly paired HK-RRs, compared to the initial random within-species pair assignments and to the final assignment predicted by the MI-IPA. The MI-IPA successfully reduces this excess pairwise MI, thus approaching the pairwise MI of the actual alignment. Same data as in (A).

MI is the actual value of the mutual information and \widehat{MI} is its naive estimate using empirical frequencies, while K is the number of independent values that can be taken by the pair of random variables considered, and M represents dataset size [32]. Estimates \widehat{MI} obtained by subsampling the initial dataset can be linearly (or polynomially, to include higher-order corrections) fitted versus $1/M$, yielding the actual MI as the intercept. However, if M is not larger than K , subleading corrections become too important, and this method becomes inaccurate. Here, K can go up to the square of the number of possible states per site, i.e. ~ 400 . Consistently, we found that polynomial fits become poor when considering datasets with less than ~ 1000 sequences. Fig. S2A presents results obtained using corrections from these fits, for datasets of at least 1000 sequences. A more sophisticated method to reliably estimating the entropies of discrete distributions was introduced in Ref. [33]. It employs a Bayesian approach and a flat prior on the entropy, and results in a correction of finite-size bias in entropy estimates. This NSB (Nemenman-Shafee-Bialek) correction is successful in undersampled cases [33, 34]. In Fig. S2B, we show results obtained on our datasets using the NSB correction. Employing either of these two corrections of finite-size effects on the MI estimates (see Fig. S2) confirms our previous conclusion. For all dataset sizes, the pairwise MI of the assignment predicted by the MI-IPA is significantly closer to that of the actual assignment than to that of the initial random assignment, thus confirming that the MI-IPA yields results with near-maximal pairwise MI.

Discussion

Here, we introduced a method based on MI to predict interacting partners among the paralogs of two protein families, starting from sequences only. Specifically, our iterative pairing algorithm (MI-IPA) finds an assignment of protein pairs that approximately maximizes the MI between the sequences of the two protein families. We demonstrated that the MI-IPA allows one to

predict interacting protein pairs with high accuracy, starting from sequence data only. High performance is obtained even in the absence of an initial training set of known interacting pairs.

In Ref. [16], us and colleagues introduced a similar iterative pairing algorithm (DCA-IPA) that approximately maximized an effective interaction energy between proteins, instead of MI. This effective interaction energy was calculated from a global statistical model, more specifically a pairwise maximum entropy model, approximately inferred from the empirical one and two-body frequencies of the sequence data. Such models, often called Direct Coupling Analysis (DCA), have been successful at predicting amino-acid pairs that are in contact in folded proteins [4, 6, 35], and have permitted prediction of the three-dimensional structure of proteins from sequence data [4, 6, 7, 13, 35, 36]. For structure prediction, these global statistical models outperform the use of MI [4, 6], and have yielded major progress in the field, despite the promising results obtained by MI-based methods implementing corrections for background MI [37]. In the specific case of interacting proteins, DCA allows one to simultaneously infer interaction partners and structural contacts between them [17].

Here, we demonstrated that the MI-IPA performs at least as well as the DCA-IPA, and often outperforms it. This result highlights a striking difference between the prediction of contacting pairs of residues, and the prediction of interacting partners. Crucially, we have achieved an accurate identification of interacting partners without the need to build a global statistical model of the sequence data highlighting effective pairwise interactions between contacting residues, which is at the root of DCA. An important motivation underlying DCA and other related approaches, such as the Bayesian network method of Ref. [3] and the sparse inverse covariance estimation of Ref. [38], is to disentangle correlations arising caused by direct interactions from indirect correlations due to a chain of couplings [2–4, 38]. While this distinction between direct and indirect correlations is crucial to infer which residues are in contact in the folded protein, its importance is probably reduced in the determination of interacting partners among paralogous proteins. Besides, MI has been extremely successful in determining “specificity residues” crucial in the interactions between HKs and their cognate RRs: these residues are those involved in interprotein residue pairs with highest MIs [39–41]. Strikingly, mutating only these few specificity residues has allowed to fully switch the specificity of HK-RR pairs [39]. Hence, our present results reinforce the conclusion of these earlier studies, showing that MI can accurately reveal the specificity of interacting protein pairs.

Finally, the MI between two sequence sites simply measures how much observing a residue at one site tells us about the other. It is entirely agnostic regarding the origin of this statistical dependence, and incorporates contributions from phylogeny [22, 23] as well as from those arising from functional selection. In addition, the latter contributions include those from structural contacts, but might also involve aspects of protein function other than structural ones, including collective correlations between residues [18–21, 42]. This stands in contrast with DCA, which was developed and optimized specifically to infer contacting residues and protein structure. The fact that the MI-IPA slightly outperforms the DCA-IPA thus constitutes a hint that sources of covariation other than those that maintain structural contacts help to identify interaction partners among paralogs. This might be due to multiple effects. In particular, global functional selection could potentially affect both interacting partners together. For instance, a functionally important mechanical deformation of the complex formed by these partners could be subject to selection, yielding collective correlations that extend in the sequences of both partners [20, 21]. Besides, interacting partners may also share more phylogenetic history than non-interacting proteins, e.g. if the genes encoding the partners tend to be duplicated and/or horizontally transferred together [43–45]. Then, phylogenetic correlations could aid the prediction of interacting partners among paralogous proteins, despite being deleterious to residue contact identification [23]. Next, it will be interesting to further investigate these various sources of covariation, both functional and phylogenetic. This should be useful for the particular problem of prediction of interacting protein pairs, as well as for the more general understanding of the sequence-function relationship in proteins.

Methods

376

Here, we first explain the different steps of the MI-IPA. Next, we describe the datasets used and the way the statistics of these alignments are computed.

377

378

Iterative pairing algorithm based on mutual information (MI-IPA)

379

Ref. [16] introduced an iterative pairing algorithm (IPA) to predict interaction partners among paralogs from two protein families. It essentially performs an approximate maximization of the average effective interaction energy between pairs of proteins comprised of one protein of family \mathcal{A} and one of family \mathcal{B} , and the effective interaction energy is calculated from a pairwise maximum entropy model (see also Ref. [17]). Here, we propose the MI-IPA, a variant of the IPA that approximately maximizes MI via S_X (see Eq. 4), instead of the effective interaction energy. Importantly, the MI-IPA does not require the construction of a global statistical model of the data, contrary to the algorithms from Refs. [16, 17]. Let us now describe each step of an iteration of the MI-IPA, after explaining how the CA is constructed for the very first iteration.

380

381

382

383

384

385

386

387

388

For simplicity, we assume that in each species, there is the same number of proteins of family \mathcal{A} and of family \mathcal{B} . If this is not the case, an injective matching strategy can be used [17], so in each species, the relevant number of proteins is the minimum of the number of proteins of family \mathcal{A} and of proteins of family \mathcal{B} .

389

390

391

392

Initialization of the CA

393

If starting from a training set of known interaction partners AB, the CA for the first iteration of the IPA is built from the pairs AB in this training set. In subsequent iterations, the training set pairs are always kept in the CA, and additional pairs with the highest confidence scores (see below) are added to the CA.

394

395

396

397

In the absence of a training set, each protein A of the dataset where we wish to predict pairings is randomly paired with a protein B from its species. All these random pairs are included in the CA for the first iteration of the MI-IPA. Hence, this initial CA contains a mixture of correct and incorrect pairs, with one correct pair per species on average. At the second iteration, the CA is built using only the $N_{\text{increment}}$ AB pairs with the highest confidence scores obtained from this first iteration.

398

399

400

401

402

403

Now that we have explained the initial construction of the CA, let us describe each step of an iteration of the MI-IPA.

404

405

1. Calculation of pointwise mutual informations (PMI)

406

Each iteration begins by the calculation of PMI scores from the CA of paired AB sequences. The empirical one- and two-site frequencies, $f_i(\alpha)$ and $f_{ij}(\alpha, \beta)$, of occurrence of amino-acid states α (or β) at each site i (or j) are computed for the CA, using a specific weighting of similar sequences, and a pseudocount correction (see below) [4, 6, 11, 35]. The PMI of each residue pair (α, β) at each pair of sites (i, j) , defined in Eq. 1, is then estimated from these frequencies as

407

408

409

410

411

$$\text{PMI}_{ij}(\alpha, \beta) = \log \left[\frac{f_{ij}(\alpha, \beta)}{f_i(\alpha)f_j(\beta)} \right]. \quad (9)$$

2. Calculation of pairing scores for all possible AB pairs

412

Having calculated PMIs on the CA, we next turn to the dataset where we wish to predict pairings. The MI-based pairing score S_{AB} of each possible AB pair within each species in the dataset is calculated by summing all the inter-protein PMIs involved, as defined in Eq. 3.

413

414

415

3. AB pair assignments and ranking by confidence score

We make one-to-one AB pairs within each species in the dataset by maximizing the sum of the scores of all pairs in this species. Considering the matrix of the scores of all possible pairs, this amounts to choosing one element per line and per column, such that the sum of all of them is maximal. This assignment problem is solved exactly and efficiently using the Hungarian algorithm (also known as the Munkres algorithm) [46–48]. Each pair is scored by a confidence score ΔS_{AB} , which is the difference between the sum of the scores of the assigned pairs in the species and that obtained by using the Hungarian algorithm again while this pair is disallowed. Once pairs are made and confidence scores are calculated in each species, all the assigned AB pairs from all species are ranked in order of decreasing confidence score.

Note that in Ref. [16], we had used an approximate greedy approach to the assignment problem, where the pair with lowest energy is selected first, and the two proteins involved are removed from further consideration, and the process is repeated until all As and Bs are paired. For the DCA-IPA of Ref. [16], the greedy algorithm performed marginally better than the Hungarian algorithm, while in the present MI-IPA, the Hungarian algorithm yields better performance (see Fig. S3). This slight difference may be explained by the fact the PMI scores are most meaningful collectively, as their sum tends to a pairwise approximation of MI (see Eq. 5).

4. Incrementation of the CA

The ranking of the AB pairs is used to pick those pairs that are included in the CA at the next iteration. Pairs with a high confidence score are more likely to be correct because there was less ambiguity in the assignment. The number of pairs in the CA is increased by $N_{\text{increment}}$ at each iteration, and the MI-IPA is run until all the As and Bs in the dataset have been paired and added to the CA. In the last iteration, all pairs assigned at the second to last iteration are included in the CA.

If starting from a training set of AB pairs, the N_{start} training pairs remain in the CA throughout. The As and Bs from all the other pairs in the CA are re-paired and re-scored at each iteration, and only re-enter the CA if their confidence score is sufficiently high. In other words, at the first iteration, the CA only contains the N_{start} training pairs. Then, for any iteration number $n > 1$, it contains these exact same N_{start} training pairs, plus the $(n - 1)N_{\text{increment}}$ assigned AB pairs that had the highest confidence scores at iteration number $n - 1$.

In the absence of a training set, all As and Bs in the dataset are paired and scored at each iteration, and all the pairs of the CA are fully re-picked at each iteration based on the confidence score. For any iteration number $n > 1$, the CA contains the $(n - 1)N_{\text{increment}}$ assigned AB pairs that had the highest confidence scores at iteration number $n - 1$.

Once the new CA is constructed, the next iteration can start.

Dataset construction

We use the HK-RR datasets described in Ref. [16]. Briefly, the complete dataset was built using the online database P2CS [49, 50], yielding a total of 23,424 HK-RR pairs (known by genome proximity) from 2102 different species. We focused on the protein domains through which HKs and RRs interact, which are the Pfam HisKA domain present in most HKs (64 amino acids) and the Pfam Response_reg domain present in all RRs (112 amino acids).

In most of the paper, and as in [16], we focused on a smaller “standard dataset” extracted from this complete dataset, both because protein families that possess as many members as the HKs and RRs are atypical, and in view of computational time constraints. This standard dataset was constructed by picking species randomly, and comprises 5064 pairs from 459 species. In our datasets, we discarded sequences from species that contain only one pair, for which pairing is unambiguous. It allows us to assess the impact of training set size (N_{start}) without including an implicit training set via these pairs, and it also enables us to address prediction in the absence of any known pairs (no training set).

While we used HK-RRs as the main benchmark for the MI-IPA, we assessed the generality of its performance by applying it to several other pairs of protein families. For these proteins, paired alignments of homologs of known *Escherichia coli* interacting protein pairs were built using a method adapted from Ref. [13], as detailed in Ref. [16].

Statistics of the concatenated alignment (CA)

Let us consider a CA of paired AB sequences. At each site $i \in \{1, \dots, L\}$, where L is the length of an AB sequence, a given concatenated sequence can feature any of 21 amino acid states α .

The raw empirical frequencies $f_i(\alpha)$ and $f_{ij}(\alpha, \beta)$, obtained by counting the sequences where given residues occur at given sites and dividing by the number M of sequences in the CA, are subject to sampling bias, due to phylogeny and to the choice of species that are sequenced [6, 35]. Hence, we use a standard correction that re-weights “neighboring” concatenated sequences with mean Hamming distance per site $< \theta$. In practice $\theta = 0.15$ was found to yield the best results, but the dependence of performance on θ is weak (see Fig. S4A). The weight associated to a given concatenated sequence a is $1/m_a$, where m_a is the number of neighbors of a within the threshold (including the sequence considered) [6, 11, 35]. This allows one to define an effective total number of sequences, $M_{\text{eff}} = \sum_{a=1}^M 1/m_a$.

We also introduce pseudocounts via a parameter Λ [4, 6, 11, 35] to avoid issues due e.g. to amino-acid pairs that never appear. Indeed, those can yield mathematical difficulties, such as diverging PMI estimates. Note that pseudocounts are widely used in DCA too [35]. The corrected one-body frequencies \tilde{f}_i read

$$\tilde{f}_i(\alpha) = \frac{\Lambda}{q} + (1 - \Lambda)f_i(\alpha), \quad (10)$$

where $q = 21$ is the number of possible states per site. Similarly, the corrected two-body frequencies \tilde{f}_{ij} read

$$\tilde{f}_{ij}(\alpha, \beta) = \frac{\Lambda}{q^2} + (1 - \Lambda)f_{ij}(\alpha, \beta) \text{ if } i \neq j, \quad (11)$$

$$\tilde{f}_{ii}(\alpha, \beta) = \frac{\Lambda}{q}\delta_{\alpha\beta} + (1 - \Lambda)f_{ii}(\alpha, \beta) = f_i(\alpha)\delta_{\alpha\beta}, \quad (12)$$

where $\delta_{\alpha\beta} = 1$ if $\alpha = \beta$ and 0 otherwise. We found that $\Lambda = 0.15$ yields the best performance of the MI-IPA, but the dependence of performance on Λ is weak (see Fig. S4B). Note that this value is lower than the typical value used in DCA ($\Lambda = 0.5$) [6, 35].

Note that our demonstration that the average pairing score S_X of an assignment tends to the sum of MIs of inter-protein residue pairs (see Eqs. 4 and 5) did not include proximity weightings and pseudocounts. In practice, at each iteration of our algorithm, assignments are made within each species separately, a scale at which it is convenient to just use the sum of pairing scores S_{AB} (see Eq. 3), and find the assignment that maximizes it. This means that the convergence to the sum of MIs is approximate when using proximity weightings and pseudocounts to calculate the frequencies in the CA and estimate PMIs.

Note also that empirical frequencies without weightings or pseudocounts were employed to study how well the MI-IPA maximizes MI (Figs. 4 and S2).

Acknowledgments

A.-F. B. thanks Ned S. Wingreen for illuminating discussions. A.-F. B. acknowledges the Aspen Center for Physics, which is supported by NSF Grant PHY-1607611, and more particularly the participants of the 2017 Aspen Working Group on “Covariance Analysis of Proteins” for stimulating discussions.

Supplementary figures

503

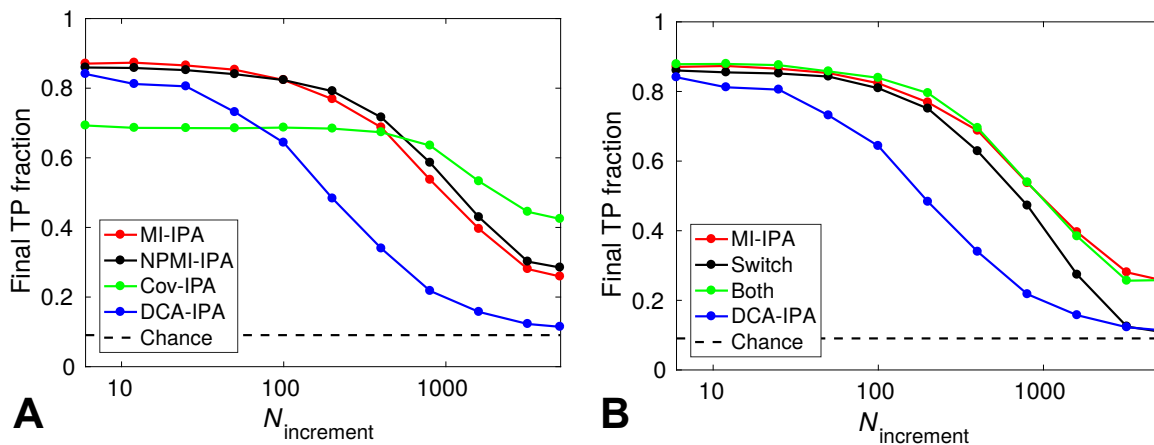


Figure S1: **IPA variants.** (A) Final TP fraction obtained versus $N_{\text{increment}}$ for the standard HK-RR dataset, starting from random within-species HK-RR pairs. In addition to the MI-IPA and the DCA-IPA (see Fig. 2A), two other variants of the IPA are presented, which score inter-protein residue pairs by the normalized PMI (NPMI-IPA) and by the covariance (Cov-IPA), respectively (see Results). Apart from this scoring difference, all particulars are the same as in the MI-IPA (see Methods). (B) Similar graph as in (A), showing two variants of the IPA that combine MI and DCA: “Switch” uses the MI-IPA for the first half of iterations, before switching to the DCA-IPA; “Both” uses both MI and DCA at each iteration and favors protein pairs that are predicted by both methods (see Results). In both panels, all results are averaged over 50 replicates employing different initial random pairings, and the dashed line represents the average TP fraction obtained for random within-species HK-RR pairings.

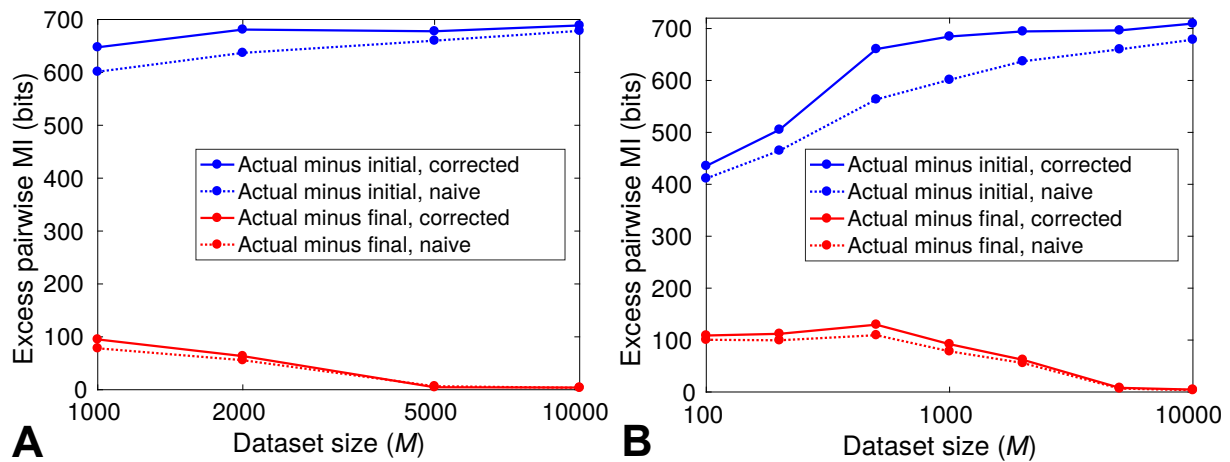


Figure S2: **Increase of pairwise MI obtained by the MI-IPA: finite-size effect corrections.** In both panels, similar results as those in Fig. 4B are reported, but two different finite-size effect corrections to the MI estimates are implemented. Excess pairwise MI of the actual set of correctly paired HK-RRs, compared to the initial random within-species pair assignments and to the final assignment predicted by the MI-IPA, are plotted versus the total number M of sequences in the dataset. (A) Finite-size correction using subsampling of each dataset and fitting versus $1/M$ (see Results). Third-degree polynomial fitting was employed. (B) NSB correction [33, 34] (see Results). In both panels, naive estimates (see Fig. 4B) are also shown for comparison. For each M , HK-RR datasets are constructed by picking species randomly from the full dataset, and results are averaged over 50 different such alignments.

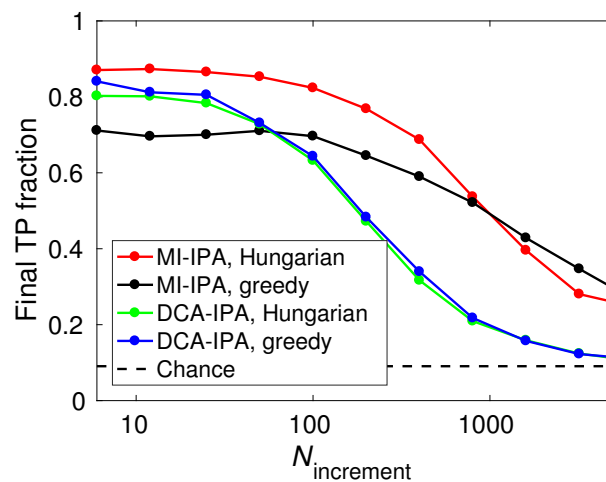


Figure S3: **IPA assignment variants.** Final TP fraction obtained versus $N_{\text{increment}}$ for the standard HK-RR dataset, starting from random within-species HK-RR pairs. For both the MI-IPA and the DCA-IPA, two variants of the protein pair assignment strategy are presented. Once pairing scores are computed, pairs are assigned within each species either using the Hungarian algorithm, which maximizes the sum of scores in the species, or using a greedy algorithm, that favors individual pairs with top scores (see Methods). In all the rest of the paper, the Hungarian algorithm is used for the MI-IPA, and the greedy one is used for the DCA-IPA, as in Ref. [16]. All results are averaged over 50 replicates employing different initial random pairings, and the dashed line represents the average TP fraction obtained for random within-species HK-RR pairings.

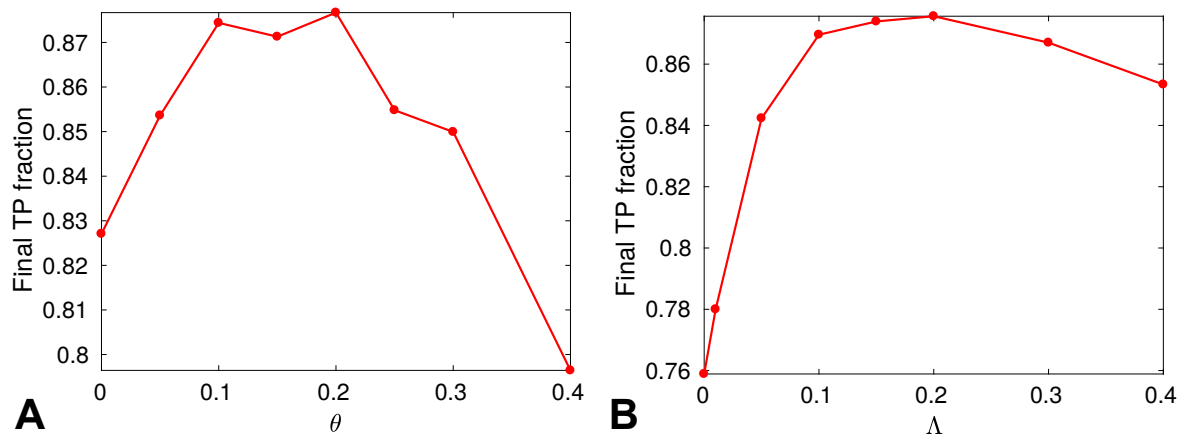


Figure S4: **Impact of the weighting of similar sequences and of the pseudocount.** (A) Final TP fraction obtained for the standard HK-RR dataset, starting from random within-species HK-RR pairs, with $N_{\text{increment}} = 6$, for various values of the threshold θ of mean Hamming distance per site below which two sequences are considered as neighbors and weighted accordingly (see Methods). Pseudocount $\Lambda = 0.15$ was used. (B) Similar graph, but varying the pseudocount Λ (see Methods), at fixed $\theta = 0.15$. In both panels, all results are averaged over 50 replicates employing different initial random pairings.

References

- [1] Rajagopala SV, Sikorski P, Kumar A, Mosca R, Vlasblom J, Arnold R, et al. The binary protein-protein interaction landscape of *Escherichia coli*. *Nat Biotechnol*. 2014 Mar;32(3):285–290.
- [2] Lapedes AS, Giraud BG, Liu L, Stormo GD. Correlated mutations in models of protein sequences: phylogenetic and structural effects. In: *Statistics in molecular biology and genetics - IMS Lecture Notes - Monograph Series*. vol. 33; 1999. p. 236–256.
- [3] Burger L, van Nimwegen E. Accurate prediction of protein-protein interactions from sequence alignments using a Bayesian method. *Mol Syst Biol*. 2008;4:165.
- [4] Weigt M, White RA, Szurmant H, Hoch JA, Hwa T. Identification of direct residue contacts in protein-protein interaction by message passing. *Proc Natl Acad Sci USA*. 2009 Jan;106(1):67–72.
- [5] Jaynes ET. *Information Theory and Statistical Mechanics*. *Phys Rev*. 1957;106(4):620–630.
- [6] Marks DS, Colwell LJ, Sheridan R, Hopf TA, Pagnani A, Zecchina R, et al. Protein 3D structure computed from evolutionary sequence variation. *PLoS ONE*. 2011;6(12):e28766.
- [7] Sulikowska JI, Morcos F, Weigt M, Hwa T, Onuchic JN. Genomics-aided structure prediction. *Proc Natl Acad Sci USA*. 2012 Jun;109(26):10340–10345.
- [8] Dwyer RS, Ricci DP, Colwell LJ, Silhavy TJ, Wingreen NS. Predicting functionally informative mutations in *Escherichia coli* BamA using evolutionary covariance analysis. *Genetics*. 2013 Oct;195(2):443–455.
- [9] Cheng RR, Morcos F, Levine H, Onuchic JN. Toward rationally redesigning bacterial two-component signaling systems using coevolutionary information. *Proc Natl Acad Sci USA*. 2014 Feb;111(5):E563–571.
- [10] Figliuzzi M, Jacquier H, Schug A, Tenaillon O, Weigt M. Coevolutionary Landscape Inference and the Context-Dependence of Mutations in Beta-Lactamase TEM-1. *Mol Biol Evol*. 2016 Jan;33(1):268–280.
- [11] Procaccini A, Lunt B, Szurmant H, Hwa T, Weigt M. Dissecting the specificity of protein-protein interaction in bacterial two-component signaling: orphans and crosstalks. *PLoS ONE*. 2011;6(5):e19729.
- [12] Baldassi C, Zamparo M, Feinauer C, Procaccini A, Zecchina R, Weigt M, et al. Fast and accurate multivariate Gaussian modeling of protein families: predicting residue contacts and protein-interaction partners. *PLoS ONE*. 2014;9(3):e92721.
- [13] Ovchinnikov S, Kamisetty H, Baker D. Robust and accurate prediction of residue-residue interactions across protein interfaces using evolutionary information. *Elife*. 2014;3:e02030. Available from: <http://gremlin.bakerlab.org/>.
- [14] Hopf TA, Scharfe CP, Rodrigues JP, Green AG, Kohlbacher O, Sander C, et al. Sequence co-evolution gives 3D contacts and structures of protein complexes. *Elife*. 2014;3:e03430.
- [15] Feinauer C, Szurmant H, Weigt M, Pagnani A. Inter-Protein Sequence Co-Evolution Predicts Known Physical Interactions in Bacterial Ribosomes and the Trp Operon. *PLoS ONE*. 2016;11(2):e0149166.
- [16] Bitbol AF, Dwyer RS, Colwell LJ, Wingreen NS. Inferring interaction partners from protein sequences. *Proc Natl Acad Sci USA*. 2016;113(43):12180–12185.

- [17] Gueudre T, Baldassi C, Zamparo M, Weigt M, Pagnani A. Simultaneous identification of specifically interacting paralogs and interprotein contacts by direct coupling analysis. *Proc Natl Acad Sci USA*. 2016 10;113(43):12186–12191.
- [18] Halabi N, Rivoire O, Leibler S, Ranganathan R. Protein sectors: evolutionary units of three-dimensional structure. *Cell*. 2009;138(4):774–786.
- [19] Rivoire O, Reynolds KA, Ranganathan R. Evolution-based functional decomposition of proteins. *PLoS Comput Biol*. 2016;12(6):e1004817.
- [20] Yan L, Ravasio R, Brito C, Wyart M. Architecture and coevolution of allosteric materials. *Proc Natl Acad Sci USA*. 2017;114(10):2526–2531.
- [21] Wang SW, Bitbol AF, Wingreen NS. Revealing evolutionary constraints on proteins through sequence analysis. Submitted;
- [22] Casari G, Sander C, Valencia A. A method to predict functional residues in proteins. *Nat Struct Biol*. 1995 Feb;2(2):171–178.
- [23] Qin C, Colwell LJ. Power law tails in phylogenetic systems. *Proc Natl Acad Sci USA*. 2018 Jan;115(4):690–695.
- [24] Fano R. *Transmission of Information: A Statistical Theory of Communications*. Cambridge, MA: The MIT Press; 1961.
- [25] Cover TM, Thomas JA. *Elements of Information Theory*. New York, NY, USA: Wiley-Interscience; 2006.
- [26] Church KW, Hanks P. Word association norms, mutual information, and lexicography. *Comput Linguist*. 1990;16(1):22–29.
- [27] Role F, Nadif M. Handling the impact of low frequency events on co-occurrence based measures of word similarity. In: Filipe J, Fred A, editors. *Proceedings of the International Conference on Knowledge Discovery and Information Retrieval (KDIR-2011)*. Scitepress; 2011. p. 218–223.
- [28] Laub MT, Goulian M. Specificity in two-component signal transduction pathways. *Annu Rev Genet*. 2007;41:121–145.
- [29] Rees DC, Johnson E, Lewinson O. ABC transporters: the power to change. *Nat Rev Mol Cell Biol*. 2009 Mar;10(3):218–227.
- [30] Ishikawa M, Tsuchiya D, Oyama T, Tsunaka Y, Morikawa K. Structural basis for channeling mechanism of a fatty acid beta-oxidation multienzyme complex. *EMBO J*. 2004 Jul;23(14):2745–2754.
- [31] Dietzel U, Kuper J, Doebbler JA, Schulte A, Truglio JJ, Leimkuhler S, et al. Mechanism of Substrate and Inhibitor Binding of *Rhodobacter capsulatus* Xanthine Dehydrogenase. *J Biol Chem*. 2009 Mar;284(13):8768–8776.
- [32] Bialek W. *Biophysics: Searching for Principles*. Princeton University Press; 2012.
- [33] Nemenman I, Shafee F, Bialek W. Entropy and inference, revisited. In: Dietterich TG, Becker S, Ghahramani Z, editors. *Advances in Neural Information Processing Systems 14*. MIT Press; 2002. p. 471–478. Available from: <http://nsb-entropy.sourceforge.net/>.
- [34] Nemenman I, Bialek W, de Ruyter van Steveninck R. Entropy and information in neural spike trains: progress on the sampling problem. *Phys Rev E*. 2004 May;69(5 Pt 2):056111. Available from: <http://nsb-entropy.sourceforge.net/>.

- [35] Morcos F, Pagnani A, Lunt B, Bertolino A, Marks DS, Sander C, et al. Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proc Natl Acad Sci USA*. 2011 Dec;108(49):E1293–1301.
- [36] Ovchinnikov S, Kinch L, Park H, Liao Y, Pei J, Kim DE, et al. Large-scale determination of previously unsolved protein structures using evolutionary information. *Elife*. 2015 Sep;4:e09248.
- [37] Dunn SD, Wahl LM, Gloor GB. Mutual information without the influence of phylogeny or entropy dramatically improves residue contact prediction. *Bioinformatics*. 2008 Feb;24(3):333–340.
- [38] Jones DT, Buchan DW, Cozzetto D, Pontil M. PSICOV: precise structural contact prediction using sparse inverse covariance estimation on large multiple sequence alignments. *Bioinformatics*. 2012 Jan;28(2):184–190.
- [39] Skerker JM, Perchuk BS, Siryaporn A, Lubin EA, Ashenberg O, Goulian M, et al. Rewiring the specificity of two-component signal transduction systems. *Cell*. 2008 Jun;133(6):1043–1054.
- [40] Podgornaia AI, Laub MT. Determinants of specificity in two-component signal transduction. *Curr Opin Microbiol*. 2013 Apr;16(2):156–162.
- [41] Podgornaia AI, Laub MT. Protein evolution. Pervasive degeneracy and epistasis in a protein-protein interface. *Science*. 2015 Feb;347(6222):673–677.
- [42] Cunningham AD, Colavin A, Huang KC, Mochly-Rosen D. Coupling between Protein Stability and Catalytic Activity Determines Pathogenicity of G6PD Variants. *Cell Rep*. 2017 03;18(11):2592–2599.
- [43] Alm E, Huang K, Arkin A. The evolution of two-component systems in bacteria reveals different strategies for niche adaptation. *PLoS Comput Biol*. 2006 Nov;2(11):e143.
- [44] Capra EJ, Laub MT. Evolution of two-component signal transduction systems. *Annu Rev Microbiol*. 2012;66:325–347.
- [45] Rowland MA, Deeds EJ. Crosstalk and the evolution of specificity in two-component signaling. *Proc Natl Acad Sci USA*. 2014 Apr;111(15):5550–5555.
- [46] Kuhn HW. The Hungarian Method for the assignment problem. *Naval Research Logistics Quarterly*. 1955;2:83–97.
- [47] Munkres J. Algorithms for the Assignment and Transportation Problems. *Journal of the Society for Industrial and Applied Mathematics*. 1957;5(1):32–38.
- [48] Buehren M. Functions for the rectangular assignment problem; 2011. Available from: <https://fr.mathworks.com/matlabcentral/fileexchange/6543-functions-for-the-rectangular-assignment-problem>.
- [49] Barakat M, Ortet P, Jourlin-Castelli C, Ansaldi M, Mejean V, Whitworth DE. P2CS: a two-component system resource for prokaryotic signal transduction research. *BMC Genomics*. 2009;10:315. Available from: <http://www.p2cs.org/>.
- [50] Ortet P, Whitworth DE, Santaella C, Achouak W, Barakat M. P2CS: updates of the prokaryotic two-component systems database. *Nucleic Acids Res*. 2015 Jan;43(Database issue):D536–541. Available from: <http://www.p2cs.org/>.